

COARSE FREQUENCY OFFSET ESTIMATION

BACKGROUND OF THE INVENTION

TECHNICAL FIELD

The invention relates to physical layer (PHY) digital signal processing for use in processors developed for wireless local area networks (LAN's), and more particularly to wireless LAN's based on orthogonal frequency division multiplexing (OFDM) in the license-free national information structure (U-NII) radio spectrum bands in the United States and generally conforming to IEEE Specification 802.11a.

DESCRIPTION OF THE PRIOR ART

Local area networks (LAN's) have traditionally been interconnected by twisted-wire pairs and shielded cables. However, there are several deficiencies of traditional LAN's, the main being restricted mobility. In contrast, a whole class of untethered computing has emerged which uses complex modulation and coding to achieve high-speed data rates. The IEEE 802.11a standard, herein "802.11a", specifies, among other things, the physical layer (PHY) entity for an orthogonal frequency division multiplexing (OFDM) system with data payload communication capabilities of 6,9,12,18,24,36,48, and 54Mb/s. The 802.11a standard specifies RF transmission in the 5.15-5.25, 5.25-5.35, and 5.725-5.825GHZ unlicensed national information structure (U-NII) bands.

Typically, the IEEE communication standards specify the transmit bit-stream in addition to performance specifications, RF emissions requirements, etc. The wireless transmission medium inherently introduces some unique impairments (not present in traditional LAN's) to the transmitted signal which must be mitigated in the remote receiver station. These impairments

include signal fading, multi-path reflections, base- and remote-unit oscillator mismatch introduced frequency offset, timing misalignment, and timing synchronization. In addition, there are RF hardware limitations such as receiver IQ imbalance and phase noise that must be mitigated as well. As such, the mitigation of such effects falls under the category of baseband digital signal processing. To assist the remote unit in mitigating these effects, a known training sequence is usually embedded into the transmit bit stream; this occurs at the expense of bandwidth. Of course, the same problems occur in the upstream direction (remote station transmitting to the base station), but it suffices to discuss the downstream digital signal processing.

In this disclosure, one such digital signal processing method, coarse frequency estimation, is outlined. This processing block digitally estimates the oscillator mismatch between the base- and remote-station and corrects for it in subsequent data demodulation. Typical voltage-controlled temperature-compensated crystal oscillators (VCTCXO) used in wireless communications have a ± 20 (parts-per-million) ppm error. At 5 GHz (5000 MHz), this translates to an error of ± 100 kHz at each end, or ± 200 kHz in combination. With OFDM modulation, a frequency error of 3% of the inter-carrier frequency spacing is the maximum tolerable frequency error.

The transmission scheme in 802.11a is bursty. This means that the receivers must digitally process the training sequence to mitigate the undesired signal impairments each time a burst commences. This means that it is desirable for the processing blocks to be as robust and computationally efficient as possible.

The quality of carrier frequency-offset estimation must be such that the relative error between actual and estimated values does not exceed three percent of the frequency spacing between consecutive sub-carriers, *e.g.* 9.375 kHz. To reach this target precision, the 802.11a PHY specification recommends that frequency offset estimation be carried out into two

successive stages, a coarse and fine frequency estimation stage. Coarse and fine estimates must be derived from the processing of the short and long preambles respectively. See, IEEE-802.11a-1999, §17.3.3. For short, these are called the "short preamble" and the "long preamble."

Frequency offset errors must be removed for a receiver to track the transmitted signal and demodulate it properly. A conventional method exists to remove such offset which involves a control loop which feeds a frequency error signal back to a VCTCXO to slowly correct the oscillator mismatch. C&S Technology (Korea) has announced a wireless-LAN modem-chip for IEEE-802.11a applications (see <http://cnstec.com/e-html/products/products-1-2-4.htm>). Such uses an automatic frequency control (AFC) clock recovery circuit to correct frequency offset errors. However, due to the relatively short time span of the training sequence and the loop bandwidth of the control loop may result in inaccurate frequency correction. The method described herein does not use AFC circuitry. Rather, it estimates the existing frequency offset and instead of correcting for it with an AFC loop in an analog fashion, it constructs a frequency correcting cisoid at a frequency that is negative to the estimated frequency offset and uses this in subsequent digital signal processing and demodulation.

SUMMARY OF THE INVENTION

Orthogonal frequency division multiplexing (OFDM) receiver embodiments of the invention demodulate quadrature amplitude modulated (QAM) signals transmitted in the 5 GHz frequency band and digitally correct for frequency offset errors in their digital signal processing (DSP) units. Modulation types include binary phase shift keying (BPSK), quadrature phase shift keying (QPSK), 16-QAM, 64-QAM, (and 256-QAM in future standard enhancements). During the short preamble, the first few constituents are deliberately skipped (the short preamble consists of 10 repetitions of a 16 sample sequence; a constituent is one 16 sample sequence). Then every 2.4 μ s, a 1.6 μ s-duration sequence is collected until three such sequences

are collected. The way these three sequences are collected ensures that they represent the same waveform. Now, in the presence of a frequency offset, these sequences are phase offset with respect to one another (the second one with respect to the first one and the third one with respect to the second one are phase offset with the same value which directly relates to the frequency offset error). Once the frequency offset error is determined the estimate is used in subsequent digital signal processing.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a functional block diagram of an OFDM radio-transceiver embodiment of the invention;

Fig. 2 is a diagram representing the sampling of the short preamble used in training sequences for the physical layer (PHY) of a wireless local area network (LAN) conforming to the IEEE-802.11a specification; and

Fig. 3 is a diagram representing the sampling of the short preamble used in training sequences for the physical layer (PHY) of a wireless local area network (LAN) conforming to the IEEE-802.11a specification according to an alternate embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

Fig. 1 illustrates a wireless local area network (LAN) embodiment of the invention, and is referred to herein by the general reference numeral 100. Such, wireless LAN is preferably based on orthogonal frequency division multiplexing (OFDM), and quadrature amplitude modulated (QAM) of signals transmitted in the license-free 5 GHz frequency band. Modulation types include binary phase shift keying (BPSK), quadrature phase shift keying (QPSK), 16-QAM, 64-QAM, (and 256-QAM in future standard enhancements). The wireless LAN 100 typically includes a wireless network 101 connected to the Internet, a PHY-transmitter 102, and a PHY-receiver 103. Such units all conform to the IEEE-802.11a specification for a

physical layer (PHY) interface in a wireless local area network which allows mobile clients. The transmitter 102 comprises a digital signal processor (DSP) 104 which implements a forward error correction (FEC) coder 106, an interleaving and mapping process 108, an inverse fast Fourier transform processor 110, and a symbol wave shaper 112. The DSP 104 outputs in-phase (I) and quadrature-phase (Q) signals that are input to an IQ modulator 114 driven by a local oscillator 116. The modulated output is sent to a mixer 118 for upconversion to the 5 GHz band. A second local oscillator 120 provides the necessary carrier frequency. A high power amplifier (HPA) 122 drives a transmitter antenna 124. A radio up-link 125 is received by the wireless network 101. In general, the transmitter 102 can be implemented with conventional methods and components.

The receiver 103 receives a radio down-link 126 that is typically transmitted in bursts. Each burst is begun with a training sequence, *e.g.* a short and long preamble. The receiver 103 includes a receiver antenna 128 followed by a low-noise amplifier (LNA) 130. A local oscillator 132 and a first mixer 134 produce an intermediate frequency (IF). An automatic gain control (AGC) amplifier 136 smoothes out signal-strength variations and drives an IQ-detector 138. A second local oscillator 140 provides the carrier necessary to derive the I and Q samples. In all presently preferred embodiments of the invention, no automatic frequency control (AFC) clock recovery is needed because any frequency offset errors are corrected in later digital processing. A receiver-DSP 142 comprises a fast Fourier transform process 144, a demapping and deinterleaving process 146, and an FEC decoder 148. The receiver-DSP 142 further includes the necessary digital logic needed for carrier frequency offset determination and correction.

According to the IEEE-802.11a PHY specification, the frequency-offset estimation error must not exceed three percent of the adjacent sub-carrier channel spacing, *e.g.* 9.375 kHz. The specification therefore recommends that frequency offset estimation be carried out into two successive stages,

e.g. a coarse estimation followed by a fine estimation. Such estimates are respectively derived from the processing of the so-called short and long preambles. This disclosure is limited to the coarse estimation of frequency offset using the short preamble. The estimation of the coarse frequency offset is the first operation to be performed once an incoming packet has been detected. In preferred embodiments, the coarse frequency offset operation precedes any other such operations as synchronization, i.e. determination of the boundary between short and long preambles; in other words, acquisition of a timing reference, timing misalignment estimation, i.e. intra-baud timing offset acquisition or equalization, i.e. determination of the transmission channel impulse response.

The QAM encoding of the transmitter signal causes symbols to be decoded in ways that are affected by the carrier frequency offset error. Presently preferred embodiments of the invention take advantage of the fact that the short preamble repeats the same sequence ten times in succession. On the transmitter side, the short preamble is generated by taking the inverse Fourier transform of a predefined 64-sample complex sequence (frequency domain) whose non-zero values could be regarded as a QPSK signal (of course, no information is embedded into it). The resulting signal (time domain) corresponds to the repetition of a basic constituent (four times) and appending this signal to itself one time and a half generates the rest of the short preamble. Implementation-wise though, no inverse Fourier transform is computed because the entire short preamble is stored in a lookup table. The same principle applies to the generation of long preamble, except that the long preamble is obtained through a BPSK-like signal and contains two and a half basic constituents). If we accept any other phenomena, such as multipath or distortion, then any phase difference observed between short preamble constituents can be attributed to the carrier frequency offset. Therefore, if three equally spaced sequences (having equal length) are extracted from the short preamble then the solution to a simultaneous equation that includes all three sequences yields their common frequency offset error. Mathematically, the problem reduces to the calculation and

subsequent processing of a covariance matrix R characterized by a pseudo-rank of one (the rank of a matrix M is the largest number of columns that constitute a linearly independent set; in our case, since three identical (but time delayed) sequences are used to calculate R , the rank is ideally one).

5

Method embodiments of the invention are based on an eigenvector decomposition of a covariance matrix. This is computed from three equally spaced 16-sample sequences extracted from the short preamble. And that is followed by a maximization of a cost function derived from the maximum eigenvector.

10

Described mathematically, let $s_1(n), s_2(n), s_3(n)$ designate the column vectors associated with the three extracted sequences mentioned above. The covariance matrix is given by:

$$R = Z^H Z$$

With:

$$Z = [s_1(n) \quad s_2(n) \quad s_3(n)]$$

$$s_1(n) = \sigma_1(n) + \eta_1(n)$$

$$s_2(n) = \sigma_2(n) + \eta_2(n)$$

$$s_3(n) = \sigma_3(n) + \eta_3(n)$$

$$\sigma_1(n) = A e^{j2\pi \frac{v}{F_s} n + j\Phi(n) + j\varphi}$$

$$\sigma_2(n) = A e^{j2\pi \frac{v}{F_s} (n+N) + j\Phi(n+N) + j\varphi} = e^{j2\pi \frac{v}{F_s} N} \sigma_1(n)$$

$$\sigma_3(n) = A e^{j2\pi \frac{v}{F_s} (n+2N) + j\Phi(n+2N) + j\varphi} = e^{j2\pi \frac{v}{F_s} 2N} \sigma_2(n) = e^{j4\pi \frac{v}{F_s} N} \sigma_1(n)$$

Expanding (1),

$$R = \begin{bmatrix} MA^2 + \gamma_{1,1}^M & MA^2 e^{j2\pi \frac{\nu}{F_s} N} + \gamma_{1,2}^M & MA^2 e^{j4\pi \frac{\nu}{F_s} N} + \gamma_{1,3}^M \\ MA^2 e^{-j2\pi \frac{\nu}{F_s} N} + \bar{\gamma}_{1,2}^M & MA^2 + \gamma_{2,2}^M & MA^2 e^{j2\pi \frac{\nu}{F_s} N} + \gamma_{2,3}^M \\ MA^2 e^{-j4\pi \frac{\nu}{F_s} N} + \bar{\gamma}_{1,3}^M & MA^2 e^{-j2\pi \frac{\nu}{F_s} N} + \bar{\gamma}_{2,3}^M & MA^2 + \gamma_{3,3}^M \end{bmatrix}$$

With:

$$\gamma_{k,l}^M = \sum_{n=1}^M \eta_k(n) \eta_l^*(n)$$

5

In the absence of noise, R is rank 1 and decomposes as follows:

$$R = MA^2 \begin{bmatrix} 1 & e^{j2\pi \frac{\nu}{F_s} N} & e^{j2\pi \frac{\nu}{F_s} N} \\ e^{-j2\pi \frac{\nu}{F_s} N} & 1 & e^{j2\pi \frac{\nu}{F_s} N} \\ e^{-j4\pi \frac{\nu}{F_s} N} & e^{-j2\pi \frac{\nu}{F_s} N} & 1 \end{bmatrix} = MA^2 \begin{bmatrix} 1 \\ e^{-j2\pi \frac{\nu}{F_s} N} \\ e^{-j4\pi \frac{\nu}{F_s} N} \end{bmatrix} \begin{bmatrix} 1 & e^{j2\pi \frac{\nu}{F_s} N} & e^{j4\pi \frac{\nu}{F_s} N} \end{bmatrix} = \lambda a^H(\nu) a(\nu)$$

Given,

$$f(\mu) = w(\mu) R w^H(\mu) = w(\mu) a^H(\nu) a(\nu) w^H(\mu) = |w(\mu) a^H(\nu)|^2 = \left| 1 + e^{j2\pi \frac{\mu-\nu}{F_s} N} + e^{j4\pi \frac{\mu-\nu}{F_s} N} \right|^2$$

15

with,

$$w(\mu) = \begin{bmatrix} 1 & e^{j2\pi \frac{\mu}{F_s} N} & e^{j4\pi \frac{\mu}{F_s} N} \end{bmatrix}$$

20

Because R is positive definite, $f(\mu) \geq 0 \forall \mu \in R$. Inspection shows that f exhibits one strong maximum (when $\mu = \nu$) and several depressed secondary peaks.

In practical implementations, the received signals are adversely affected by additive white Gaussian noise and multi-path interference (not to mention co-channel interference). The determining of R's maximum eigenvector is a first step. A simple and cost-effective iterative method embodiment of the invention can quickly converge to the solution, e.g. within two iterations typically for a pseudo rank one matrix.

For example, expressing R and any vector $x \in C^K$ using the eigenvector basis,

$$R = \sum_{n=1}^K \lambda_n a_n^H a_n$$

$$x = \sum_{m=1}^K \gamma_m a_m$$

The vector matrix products are computed by,

$$x_1 = x_0 R = \sum_{m=1}^K \gamma_m a_m \sum_{n=1}^K \lambda_n a_n^H a_n = \sum_{m=1}^K \sum_{n=1}^K \gamma_m \lambda_n (a_m a_n^H) a_n = \sum_{m,n=1}^K \gamma_m \lambda_n \delta_{m,n} a_m = \sum_{m=1}^K \gamma_m \lambda_m a_m$$

$$x_2 = x_1 R = \sum_{m=1}^K \gamma_m \lambda_m a_m \sum_{n=1}^K \lambda_n a_n^H a_n = \sum_{m=1}^K \sum_{n=1}^K \gamma_m \lambda_m \lambda_n (a_m a_n^H) a_n = \sum_{m,n=1}^K \gamma_m \lambda_m \lambda_n \delta_{m,n} a_m = \sum_{m=1}^K \gamma_m \lambda_m^2 a_m$$

This iterative operation is repeated up to x_k ,

$$x_k = x_{k-1} R = \sum_{m=1}^K \gamma_m \lambda_m^{k-1} a_m \sum_{n=1}^K \lambda_n a_n^H a_n = \sum_{m=1}^K \sum_{n=1}^K \gamma_m \lambda_m^{k-1} \lambda_n (a_m a_n^H) a_n = \sum_{m,n=1}^K \gamma_m \lambda_m^{k-1} \lambda_n \delta_{m,n} a_m = \sum_{m=1}^K \gamma_m \lambda_m^k a_m$$

For the sake of simplicity, R is assumed to have pseudo rank one. This means that the spectrum of eigenvalues is such that $\lambda_1 = \lambda_1 \gg \lambda_2, \lambda_3, \dots, \lambda_K$.

Then x_k rapidly converges to $\lambda_1^k a_1$ as k increases.

Recall that the maximum eigenvector has the following format:

$$w(\mu) = \begin{bmatrix} 1 & e^{j2\pi \frac{\mu}{F_s} N} & e^{j4\pi \frac{\mu}{F_s} N} \end{bmatrix}$$

The frequency offset is calculated by maximizing an objective function. The objective function, $f(\mu)$ is computed on a grid of μ , where μ ranges from – 208.33kHz to 208.33kHz. The objective function is defined to be:

$$\bar{q}(v) = \begin{bmatrix} 1 \\ e^{j2\pi \frac{v}{F_s} N} \\ e^{j4\pi \frac{v}{F_s} N} \end{bmatrix}$$

$$f(u) = \|\bar{w}(\mu)\bar{q}(v)\|^2$$

Clearly, when $\mu = v$, then a maximum of $f(\mu)$ occurs. The resolution of this search grid depends on the requirements. For the presently preferred application, the required resolution is approximately +/- 10kHz. If this search grid is split into 64 equal regions, the final resolution in theory is half of the bin width (=416.66kHz/64) or 3.255kHz, which is more than sufficient. However, due to noise and other imperfections, it has been seen that partitioning the 416.66kHz search region into 64 sub-regions does not provide sufficient resolution to determine the peak to within the required resolution. As such, a method is proposed which achieves the resolution with minimum extra computation. It is described as follows:

The pre-stored steering vectors, or $q(\mu)$, are defined where μ is defined as 0:208.33kHz in 64 equal increments. The sign of the imaginary part of the second component of the computed dominant eigenvector is analyzed. This value is $\exp(j*2*\pi*(\mu/F_s)*N)$. This is enough information to determine whether the frequency offset is positive or negative. If the sign of the imaginary part of the second component of the computed dominant eigenvector is negative, the pre-stored steering vectors are conjugated. The objective function is then computed as normal and the estimated frequency is estimated and then, finally, the sign is swapped depending on the sign of the imaginary part of the second component of the computed dominant eigenvector. One advantage of this is that the search resolution is now

effectively double that of approach before and there is enough margin built in to account for peak shifting due to imperfections. This extra step, which takes minimal cycles to implement, completely eliminates the need for a fine-frequency estimation stage which reduces MIPS. Clearly, to achieve the same frequency resolution, one could double the resolution on the grid to 128, but this would require substantially more MIPS.

Fig. 2 represents a short preamble 200 used in training sequences for the physical layer (PHY) of a wireless local area network (LAN) conforming to the IEEE-802.11a specification. Such is sampled three times after a short delay time 202 from a packet transmission start 204. A first measurement 206 typically spans thirty-two sample times from a start time 208 to an end time 210. A first period delay 212 is typically forty-eight sample times until a next start time 214. A second measurement 216 is collected thirty-two more sample times until an end time 218. Similarly, a second period delay 220 is typically forty-eight sample times until a next start time 222. A third measurement 224 is collected thirty-two more sample times until an end time 226. In the absence of any noise and distortions, measurements 206, 216, and 224 differ only by their relative phase offset (the frequency offset information is embedded into these phase offsets as explicitly showed in formula (2) above) if a set of symbols 231-240 are all ten the same. All presently preferred embodiments of the invention rely on this observation and use digital signal processing methods to correct the common frequency offset that has been computed in a wireless local area network.

Alternate Embodiment for Coarse Frequency Offset Estimation

In an alternate embodiment, the coarse frequency offset is based on Capon's minimum variance estimator. This alternate embodiment improves performance when there are one or more co-channel interferers. This embodiment differs from the preferred embodiment in that the cost function used to estimate the coarse frequency offset is directly derived from Capon's estimator. Note though that the way the signal is extracted from the

short preamble stays exactly the same as in the preferred embodiment, *i.e.* one collects three equally spaces (48 samples) equal-length sequences, each comprising 32 samples in the current implementation and not 16 as actually shown.

5

Let $s_1(n), s_2(n), s_3(n)$ designate the column vectors associated with the three extracted sequences mentioned above. The covariance matrix is given by:

$$R = Z^H Z$$

10

With:

$$Z = \begin{bmatrix} s_1(n) & s_2(n) & s_3(n) \end{bmatrix}$$

$$s_1(n) = \sigma_1(n) + \eta_1(n)$$

$$s_2(n) = \sigma_2(n) + \eta_2(n)$$

$$s_3(n) = \sigma_3(n) + \eta_3(n)$$

$$\sigma_1(n) = A e^{j2\pi \frac{v}{F_s} n + j\Phi(n) + j\varphi}$$

$$\sigma_2(n) = A e^{j2\pi \frac{v}{F_s} (n+N) + j\Phi(n+N) + j\varphi} = e^{j2\pi \frac{v}{F_s} N} \sigma_1(n)$$

$$\sigma_3(n) = A e^{j2\pi \frac{v}{F_s} (n+2N) + j\Phi(n+2N) + j\varphi} = e^{j2\pi \frac{v}{F_s} 2N} \sigma_2(n) = e^{j4\pi \frac{v}{F_s} N} \sigma_1(n)$$

The embodiment is as follows:

15

in the absence of any noise and interferers, it is easy to see from the formulas above that the frequency-offset information is embedded into the phase of σ_2 and σ_3 in the form of a relative phase offset ($2\pi \frac{v}{F_s} N$ and $4\pi \frac{v}{F_s} N$ respectively) relatively to σ_1 .

20

Now, if one or more interferers are present (it is assumed that the S over I ratio is greater than, *e.g.* 10 dB), signal and interferers respective short preambles may largely overlap time-wise. Moreover, the interferers are very likely to be characterized by a different carrier offset. If the task is to combine the extracted

sequences s_1 , s_2 and s_3 so as to best estimate σ_1 (idealized signal of interest) then we would have to find a way to eliminate the unwanted contribution of interferers through appropriate weighting. This could be achieved through the use of Capon's minimum variance estimator. This estimator is well known in array processing. In fact, because of the particular way we collect s_1 , s_2 and s_3 , the problem of estimating the frequency offset is similar to that of estimating the direction of arrival (DOA) of a signal impinging on a 3-antenna array (outputting three signals). Once the DOA has been determined, Capon's estimator allows the derivation of the transmit weights that minimize the total power output by the array. In particular, nulls are steered towards interferers while constraining the gain of the array in the direction of the signal of interest to be constant. The receive weights are simply deduced from the transmit weights by taking the conjugate of the latter. Now, since our task is simply to estimate the frequency offset *i.e.* the equivalent of the DOA in array processing. We do not seek here to combine s_1 , s_2 and s_3 into one single signal, we just need one by-product of Capon's estimator, namely the distribution of the received power versus the DOA, *i.e.* the frequency offset in our case.

$$f(\mu) = \frac{1}{g(\mu)} = \frac{1}{w(\mu) R^{-1} w^H(\mu)}$$

Where:

$$w(\mu) = \begin{bmatrix} 1 & e^{j2\pi\frac{\mu}{F_s}N} & e^{j4\pi\frac{\mu}{F_s}N} \end{bmatrix}$$

is a steering vector. Hence, the value of the frequency offset, *i.e.* μ is simply derived through the minimization of $g(\mu)$ (our new cost function).

The minimization of $g(\mu)$ is conducted in a way very similar to the way the cost function is maximized in the preferred embodiment, *i.e.* by spanning a

grid of μ values through the use of an array of predefined steering vectors. The only difference lies in the fact that R^{-1} must be computed. Remember that there is no need to invert R in the preferred embodiment. Whereas in the preferred embodiment the computational complexity of the cost function was greatly reduced through the explicit computation of the maximum eigenvector (intermediate step), it is out of question to proceed likewise here for the simple reason that the pseudo rank of R^{-1} is assumed to be greater or equal than 2 as we assume that one or more interferers are present.

Performance-wise, this alternate estimator performs much better than the preferred one when one or more co-channel interferers are present (its complexity is much higher though). Nevertheless, the latter shall be preferred in the absence of interferers.

Although the invention is described herein with reference to the preferred embodiment, one skilled in the art will readily appreciate that other applications may be substituted for those set forth herein without departing from the spirit and scope of the present invention. Accordingly, the invention should only be limited by the Claims included below.